UCRL-MI-142775

ASCI ASAP

Procurement and Materiel
(University of California/DOE)
Contract W-7405-ENG-48
Department of Energy
Defense Programs

Request for Expressions of Interest
Level 2 Strategic Investigations
March 9, 2001

<u>Responses Required by April 9, 2001</u>

Dear Colleague:

Lawrence Livermore National Laboratory (LLNL), Los Alamos National Laboratory (LANL), and Sandia National Laboratories (SNL), hereinafter called the DP Laboratories, with funding provided by the U.S. Department of Energy (DOE) Defense Programs (DP), have implemented the Accelerated Strategic Computing Initiative (ASCI), which is a component of the DOE Stockpile Stewardship and Management Program. The vision for ASCI includes participation by universities to conduct strategic research as further described by the "ASCI Academic Strategic Alliances Program (ASAP) Statement" and other Internet-accessible information at the ASCI ASAP web site **(http://www.llnl.gov/asci-alliances).**

Level 2 Strategic Investigations seek to solve ASCI-related problems in specific areas by engaging individual university departments to conduct research on difficult problems in areas such as computational science and mathematics, computer science, and applications such as computational physics and mechanics or material science. A total of thirteen 3-year ASCI ASAP Level 2 Strategic Investigation projects were initiated in 1998 and 1999. These projects are now in their third and final year.

The ASAP is interested in continuing the Level 2 Strategic Investigations Program. The DP Laboratories are issuing this Request for Expressions of Interest as an invitation for your institution to express interest in participating in the next phase of the Level 2 Strategic Investigations Program by submitting White Paper(s) that describe proposed research into one or more of the problems areas identified below. The LLNL Procurement & Materiel Department is conducting this solicitation for the DP Laboratories.

The evaluation of White Papers will be done by individuals from the DP Laboratories and from DOE in accordance with the selection criteria defined elsewhere in this Request. Selected universities will receive a formal Request for Proposal (RFP) from LLNL. Proposals will be accepted only from universities that receive an RFP from the ASAP; unsolicited proposals will be returned. The ASAP anticipates that RFPs will be issued by **April 23, 2001,** with proposals due by **May 21, 2001**. The ASAP expects to have all awards in place by **October 15, 2001**.

Resulting Subcontracts may have 2- to 3-year terms and will be funded incrementally on a year-to-year basis. Overall funding for the continuation of the Level 2 program is expected to remain at $4M per year. The level of funding for a single award is expected to average $100K to $400K per year for focused, single-issue investigations. A limited number of awards (no more than two) averaging $700K to $1M per year may be available for broader, multiple-issue investigations, depending on the sophistication of the collaboration and the near-term relevance to ASCI problems. Awards are restricted to U.S. universities, although joint participation with industry, government laboratories, and other organizations is encouraged. Institutions currently working under subcontracts for the Level 1 Strategic Alliances Centers program are not eligible for selection under this Request.

ASCI-RELATED PROBLEM AREAS

Scientific Data Management, Visualization, and Discovery

Points of Contact:
Bob Tomlinson, bob@lanl.gov
Terri Quinn, quinn1@llnl.gov
Dino Pavlakos, cjpavla@sandia.gov

The ASCI scientific data management and visualization program is interested in research that will improve the effectiveness of users' analysis and exploration of extremely large and complex simulation data sets and the meshes used in these computations. Topics of interest range from interactive and automatic data exploration to improved access to archival storage systems. The three high-level focus areas can be described as: (1) Data Management, (2) Data Mining, and (3) Visualization. Systems that combine aspects of all three areas and focus on both spatial and temporal aspects of mesh data are encouraged.

Research Areas

- Interactive visualization and navigation for extremely large data sets
- Innovative visualization techniques, particularly for complicated data with many fields
- User interfaces for large tiled displays
- Automatic and semiautomatic scalable pattern recognition algorithms
- Feature characterization, identification, extraction, and tracking
- Dimension reduction techniques such as PCA and ICA
- Comparative analysis methods, both automatic and interactive

- Analysis and presentation of error, uncertainty, and other meta-information
- Data integration and data fusion of massive, diverse data collections
- Database-like querying across collections of scientific data sets and associated meta-data
- Efficient and novel access to large collections of large scientific data sets

Examples of strategies for supporting interactive exploration of extremely large data sets include multiresolution data representations and employing view-dependent and progressive rendering techniques. Automatic and semiautomatic preprocessing may be used to identify features of interest, to reduce the geometry required for visualization, and to extract structures or otherwise characterize an entire data set to assist the user in targeting regions of potential interest.

Data associated with a single project is likely to include many simulations, ranging in size from modest to extremely large, and may also include experimentally observed data. Meta-data describing these data sets is created and can be exploited in a variety of ways. Management and support for locating, extracting, deriving, and using relevant data from this large scientific data archive are required. In addition, providing the level of performance required to support analysis of these large volumes of data is a difficult challenge.

Substantial progress has been made in meeting these needs, and projects in these areas are in progress at the DP Laboratories. Proposed projects that integrate or coordinate with current tools and environments and those that leverage current and previous work will be most valuable. Plans for collaborating with DP Laboratory projects or other methods of transferring technology to the DP Laboratories are very desirable, particularly if the utility is clearly demonstrated in the context of extreme data set size and if pragmatics permit deployment in many offices.

Simulation Development Environments, Including Performance Tools

Points of Contact:
Jeffrey Brown, jeffb@lanl.gov
Neil Pundit, ndpundi@sandia.gov
John May, johnmay@llnl.gov

The development environment software for the future ASCI platforms presents many challenges that can benefit from academic research and development. It is expected that future ASCI platforms will involve use of NUMA SMP cluster architectures with well beyond 10,000 processors. ASCI production applications are in excess of 100,000 lines of code and will be required to run effectively on up to half of the full ASCI platform. It is expected that most ASCI applications development will continue to emphasize portability—of the languages, of the run-time system interfaces, and of the key development tools. To help accelerate this development environment, the ASAP is especially interested in research that will emphasize multiplatform solutions in the following areas.

Research Areas

- Performance Tools Infrastructure
  A fully scalable, multiplatform instrumentation system that can be configured or extended to develop multiple kinds of performance measurements and user interfaces.

- Performance and Scalability Analysis
  Techniques for understanding and improving application performance, including relationship to the memory hierarchy and interactions with the system environment. These may include visualization, statistics, run-time optimization, or automation, and should also include mechanisms to suggest what changes to the application or environment could improve performance. The approach proposed should be applicable to full ASCI systems.
- Parallel Models
  Extension of the environment for better NUMA utilization and dynamic parallelism.
- Code Correctness Tools
  Scalable techniques for verification of parallel (shared and distributed memory) codes and debugging of optimized code.

Meshing Generation and Management

Points of Contact:
Gerald Jungman, jungman@lanl.gov
Dave Hardin, hardin3@llnl.gov
Rob Leland, leland@sandia.gov

Mesh generation is a critical enabling technology for scientific simulations, particularly those involving complex engineering geometry. Mesh generation can require the investment of large amounts of human time and is often the dominant factor in determining the calendar time required to execute an engineering design iteration. While it might typically require a few days or a week of run time to solve an ASCI class problem on one of the supercomputing platforms, it often requires months of effort to develop the discretization supporting this simulation. Hence, realizing the full potential of massively parallel physics and engineering simulations requires that we greatly reduce the time required to mesh.

Currently, unstructured tetrahedral meshes are the only boundary-conforming meshes that can be generated nearly automatically for a previously unmeshed arbitrary geometric configuration realistic complexity. Unfortunately, many simulation codes do not accept tetrahedral meshes, and in some cases the results are questionable for those that do

The alternate, nearly automatic meshing method is to generate Cartesian Eulerian or Adaptive Mesh Refinement (AMR) meshes, superimpose the geometry on the mesh, and compute what volume fraction of each mesh element is occupied by each material or CAD part. This approach is relatively easy on the user, but can typically result in much less efficient use of the supercomputing resource, as measured by compute time, to achieve a given accuracy in the simulation. Unfortunately, Eulerian/AMR computations are not always satisfactory.

Overset (or Chimera) meshes can also reduce the human complexity of meshing, but the necessary interpolation between meshes that they require may erode the quality of the physics simulation to the point that it is unsatisfactory.

There remains a great need to accelerate the production of high-quality, boundary-conforming meshes, especially block-structured or unstructured hexahedral meshes. New, more general, automatic, and robust meshing algorithms and software implementations are essential to allow progress on solutions to this problem.

In addition to the problems unique to reducing the user effort and calendar time required to produce these classes of hexahedral meshes, there are also requirements that are common to reducing the time and effort to produce nearly all types of boundary-conforming meshes.

These common problems include exporting CAD geometry to the meshing software, simplifying (defeaturing) the CAD geometry, and healing frequently “dirty” CAD geometry. In addition, geometry, such as the space between assembled parts, which may be filled by air, gaskets, glue, etc., is frequently missing from the CAD description. Furthermore, assembled geometry may be different from the individual parts described by typical manufacturing CAD geometry (e.g., parts are distorted or modified by being welded or press-fit together). There is a strong need for theory and algorithmic work on quality optimization of previously generated meshes. Techniques for continuously revising meshes in response to changes in the underlying geometry are needed to support design optimization studies. The ASCI program requires improvements in all of the following meshing areas.

Research Areas

- CAD geometry
    - Export, simplification, and repair
    - Definition and integration of missing geometry
    - Post-CAD modification of assembled geometry
- Faster (in calendar time), robust, boundary-conforming meshing for arbitrary geometry
- Unstructured hexahedral meshing (the most widely acceptable unstructured mesh topology within ASCI)
- Block-structured meshing (which is also acceptable to unstructured hexahedral mesh codes)
- Unstructured, tetrahedral, and hybrid (hex/prism/pyramid/tet) meshes and even arbitrary connectivity meshes. (These are acceptable to a few ASCI codes, and robust, fully automatic versions of these methods are useful for representing geometry when computing volume fractions for Eulerian/AMR meshes.)
- Extrapolation of previously meshed similar geometry to accelerate meshing subsequent design iterations

Applications of massively parallel computing to solving any of the above problems are also sought. ASCI has needs for meshes of up to a billion elements or more. Furthermore, as a general principle, ASCI wishes to substitute machine time for human and calendar time.

Scalable Numerical Solvers and Algorithms

Points of Contact:
Joel Dendy, jed@lanl.gov
Gary Doolen, gdd@lanl.gov
Rob Falgout, rfalgout@llnl.gov
David Womble, dewombl@sandia.gov

One of the distinguishing features of the algorithmic work within the ASCI program is the issue of scale, both in terms of the problem size and in terms of the number of processors. For example, simulations of physical systems often involve behavior at both microscales and macroscales in time and space. If not handled properly, the result can be low accuracy in the simulations or severe limits on the size of the problem that can be solved. In short, the solver becomes the limiting factor in the development of modeling and simulation capabilities.

Many existing algorithms and research efforts are focused on the issue of scalability. For example, multigrid methods use a hierarchy of grids to attempt to solve linear systems of equations in time proportional to the number of unknowns. In addition to the solution of linear systems of equations, algorithmic research is required in such areas as nonlinear solvers, eigensolvers, adaptive methods, time integrators, and optimization. The ASAP is interested in research in the following areas.

Research Areas

- Multigrid methods
- Scalable, parallel solvers for large linear systems of equations, including multilevel methods and preconditioners, stable and robust algebraic preconditioners for iterative solvers, block iterative methods for multiple right-hand sides, and efficient parallel solver implementations for machines with complex memory hierarchies
- Specialized solvers for problems arising, for example, in electrical modeling, structural dynamics, reacting flows, and radiation transport applications
- Solvers for nonlinear systems of equations addressing such questions as convergence rates and radii of convergence, linearization schemes, and globalization methods
- Algorithms for efficient parallel computation of eigenvalues/eigenvectors of large linear systems
- Algorithms that couple linear solvers, nonlinear solvers, and time integration, including for example, adaptive convergence criteria for linear subproblems, operator splitting, and predictor-corrector methods
- Algorithms for optimization and design, including, for example, surrogate modeling, SQP methods, pattern search methods, and mixed integer programming
- Adaptive meshing, including error estimators and indicators, mesh refinement, and algorithms that couple solver and mesh adaptivity
- Algorithms for uncertainty estimation in large-scale simulations, including, for example, sensitivity analysis, sampling methods, and the solution of stochastic partial differential equations

Scalable Parallel I/O and File Systems

Points of Contact:
Gary Grider, ggrider@lanl.gov
Terry Jones, trj@llnl.gov
Lee Ward, lee@sandia.gov

Major problems facing the ASCI program today are I/O and file system issues related to storage and retrieval of terabytes to petabytes of data. This is exacerbated by the need for expensive, special purpose in-house solutions to satisfy the extreme bandwidth demands required by 10–30 TeraOPS computations. Within 5 years, ASCI scientists will require scalable I/O and file systems orders of magnitude improved in capacity, bandwidth, latency, and transaction rate. To maintain a balance with future 100 TeraOPS ASCI platforms, file systems in 2004–2005 will be expected to deliver sustained throughput of 100 or more gigabytes per second, together with security and easy access. Current market forces are not necessarily driving industry sufficiently to meet these requirements within this time frame. Reaching ASCI I/O goals will require accelerating and influencing new advances in academic research efforts as well as commercial components. It is critically important that ASCI accelerate the development and integration of new scalable, parallel I/O techniques and file system components over the next 3 to 5 years.

New storage system architectures, such as Network-Attached-Storage (NAS), Storage Area Networks (SANs), Internet SCSI (iSCSI), and Object-based Storage Devices (OSD) are rapidly emerging and gaining wider acceptance. Some of these approaches may help provide uniform accessibility, heterogeneous data sharing, and efficient layout and space allocation, but they do not necessarily solve ASCI-specific problems of high bandwidth performance and security requirements.

Another difficulty that arises in using new I/O techniques and file systems is the need for standardized programming interfaces. Old serial I/O interfaces are inadequate to use parallelism effectively. Therefore, new inherently parallel interfaces, such as MPI-IO, have been developed and provided to sit atop vendor-proprietary or research file systems and their underlying storage devices. Critical issues that must be addressed are widely distributed users, heterogeneous hosts, scalable performance, scalable access, scalable administration, and ASCI need-to-know security requirements. In addition, all of the usual requirements of a file system are assumed to remain in place, such as POSIX compliance, normal locking capabilities, persistence, integrity, and stability.

Different approaches are now under investigation to provide high-speed access to huge files stored across a large number of disk drives. These drives will likely be connected to many different hosts, such as ASCI platforms, visualization and data mining servers, capacity compute engines, and archival storage servers. Aggregating individual bandwidth from hundreds or thousands of commodity disk drives connected via high-speed networking (e.g., Gigabit Ethernet or InfiniBand) may help achieve the multiple GB/s performance demanded by future ASCI applications.

ASCI continues to work closely and successfully with industry and academic collaborators to improve performance and reliability of high-performance file systems (such as GPFS), archives (such as HPSS), and high-level I/O libraries (such as MPI-IO). However, many critical issues remain unsolved, as noted above.

Research Areas

- The development and integration of new scalable, parallel I/O techniques and file system components over the next 3 to 5 years
- High bandwidth performance
- Security requirements
- I/O techniques and file systems that address widely distributed users, heterogeneous hosts, scalable performance, scalable access, and scalable administration
- Aggregating individual bandwidth from hundreds or thousands of commodity disk drives connected via high-speed networking

Advanced Systems Architecture

Points of Contact:
Adolfy Hoisie, hoisie@lanl.gov
Karl-Heinz Winkler, khw@lanl.gov
Mark Seager, seager@llnl.gov
Jim Tompkins, jltomki@sandia.gov

The ASCI Program has pushed computer development from the first ASCI machine, ASCI Red at ~1.8 TeraOPS, to the current ASCI White machine at 12.3 TeraOPS and the 30 TeraOPS ASCI Q machine that will be fully installed in 2002. Experience with these machines has demonstrated the importance of parallel efficiency and reliability to achieving ASCI goals. There has also been substantial growth in floor space, power, and cooling requirements from machine to machine.

The ASCI Program needs new, general purpose, massively parallel, scientific computers with hundreds to thousands of TeraOPS of performance that focus on five fundamental challenges: power, floor space, cooling, interconnect scaling, and distance to memory (single CPU delivered performance). Solving these challenges will require fundamentally new computer architectures.

In addition, system reliability, availability, and serviceability are major issues of concern. The ASCI Program is interested in research for new computer architectures that are targeted at ASCI applications, that address these issues, that are based primarily on commercial, off-the-shelf technology, and that follow Moore's Law for cost/performance scaling. Successful research is not expected to provide prototypes, but should demonstrate scalability by some means, e.g., simulation.

Research Areas

- Power
- Floor space
- Cooling
- Interconnect scaling
- Distance to memory
- System reliability

- Availability
- Serviceability

Uncertainty Quantification in Computational Simulation

Points of Contact:
Cynthia Nitta, nitta1@llnl.gov
Martin Pilch, mpilch@sandia.gov
Ken Koch, krk@lanl.gov

Uncertainty Quantification (UQ) for Computational Simulation within the ASCI Verification and Validation Program is a broad topic, covering uncertainty in the software algorithms, in their implementation, in the physical and material models represented, in the input physical data, in the simulated output parameters, and in the data to which the simulation results are compared. Several statistical methods exist that can be used to characterize some of these sources of uncertainty and evaluate their impact on modeling. Traditional techniques for estimating uncertainty depend largely on an assumption that the ensemble being evaluated is composed of independent elements. This is typically not the case for most of the sources of uncertainty that are important for ASCI-scale computational simulations and their application. Qualitative judgments have historically been relied upon to provide a "feel" for which parameters were most important or uncertain. With increased reliance on computational simulation rather than on expensive experiments for high-consequence applications, more quantitative approaches are required. Topics of interest include identification, development, and evaluation of potential techniques for uncertainty quantification in the following areas of interest, with the goal of applying some of these techniques for quantifying uncertainty and sensitivities in ASCI-level simulations.

Research Areas

- UQ and sensitivity analysis methods for simulation software applied to multiscale, strongly coupled, nonlinear systems
- UQ strategies that tie together the many disparate sources of uncertainty in computational simulation, some of which are represented by very large sets of information and others that are represented by very sparse sets of information
- UQ methods that provide the maximum information for minimal intrusion into existing software

Some techniques that have been used in quantification of uncertainty have actually been sensitivity studies where, for example, a single input variable or parameter setting has been varied to determine the quantitative effect of the change on an output parameter of interest. This brute force method, even with the advent of the ASCI Program, is prohibitively expensive because a large number of such parameters are routinely varied. Where some type of uncertainty was estimated in these sensitivity studies, the sample sets of both experimental data or simulations have been very small. The quantity and quality of the data has been limited; rarely have two real experiments been truly the same. Further, there are not many cases where important experiments were repeated for purposes of gathering such statistical information.

In the ASCI Program, very large simulation codes are being developed in several application areas. An automated approach to uncertainty quantification and sensitivity analysis is highly desirable. Much of the coding is already in place, so minimally intrusive uncertainty processing software would also be highly desirable. Ideally, techniques should be implemented such that a numerical processing module could sit outside the simulation code, receiving and analyzing relevant information from a simulation as it was running. The efficiency of any potential technique should be evaluated and demonstrated to be greater than that of other brute force methods that return equivalent information. Of course, techniques that provide information unavailable by other means would be highly valued.

The ASAP will entertain research applications in the following areas as long as they are tightly linked to one or more of the areas discussed above.

Computational Condensed Matter Physics and Materials Science

Points of Contact:
Robert Albers, rca@lanl.gov
Christian Mailhiot, mailhiot1@llnl.gov
Eugene Hertel, esherte@sandia.gov

Proposals are solicited in areas focused on algorithm development, applied mathematics, software tools, and visualization specific to advanced materials simulations and whose objective is to fully exploit scalable, terascale, massively parallel platforms for high-performance materials modeling applications of relevance to ASCI-related problems.

A goal of the ASCI program is to advance the development of physics-based models to predict the dynamic response of materials under extreme conditions of pressure, temperature, high strain, and high strain rates. Materials of interest include correlated *f*-electron materials, transition metals, hot dense hydrogen, high explosives and energetic materials, polymers, and ceramics. Materials properties of relevance to ASCI mission needs include thermodynamic properties (equation of state, phase diagram, pressure- and temperature-induced structural phase transformations, etc.), constitutive properties (strength, plastic flow, etc.), failure and fracture, materials aging, and materials processing.

Research Areas

- Comprehensive materials modeling and simulation programs are being developed to predict materials properties and response at various length scales—from the quantum level to the continuum. Examples of advanced materials simulations where improved numerical algorithms, computational mathematics, visualization, and software tools are of interest include:
    - ab initio electronic structure methods, computational quantum chemistry methods, quantum molecular dynamics, and quantum Monte Carlo methods for simulations at the quantum level;
    - molecular dynamics and kinetic Monte Carlo methods for simulations at the atomic level;
    - three-dimensional dislocation dynamic simulations of microstructure evolution for descriptions of materials response at the microscale and mesoscale;
    - continuum finite-element methods for simulation of materials response at the macroscale;
    - physically based models that couple material mechanics, damage evolution, and chemical energy release in energetic materials; or

- physically based models that include species chemistry of energy release in energetic materials

Computational Physics and Computational Mechanics

Points of Contact:
Warner Miller, wam@lanl.gov
Paul Nowak, pnowak@llnl.gov
Hal Morgan, hsmorga@sandia.gov

Computational Physics

The goal is to develop strategies that will produce the desired physics solutions to problems in radiation transport, particle transport, and hydrodynamics in the lowest possible wall-clock time on computational platforms of interest to the ASCI program. The best known serial algorithms (or parallel, if they exist) should be used as starting points to devise algorithms and coding strategies that yield the best possible performance on modern, parallel, cache-based systems. Attention should be focused on devising optimal algorithms for a particular physics application and not on general-purpose algorithms.

The primary goal is to build physics algorithms that obtain a desired level of accuracy with an economy of effort (i.e., unknowns and iterations); hence, interest is high in adaptive algorithms. For hydrodynamics, this includes unstructured-grid arbitrary Lagrangian–Eulerian (ALE) methods and AMR techniques. For both ALE and AMR, coupling to sub-grid scale models for material mixing and turbulence (with comparisons to experiment) is of high interest. Other schemes such as high-order Lagrangian, ALE, and smooth-particle methods (SPH) are also desirable.

In the areas of radiation and particle transport, local refinement and adaptation in space and angle are highly desirable. In particular, novel and adaptive angular differencing techniques that accurately model discontinuities in particle direction are of high interest. Given the implicit nature of transport algorithms, improvements to iterative solution techniques are crucial. Finally, spatial discretization methods that maintain a high order of accuracy on unstructured and/or distorted grids are significant.

An equally important goal is to build algorithms that are highly efficient on parallel, cache-based computers. Attention should be given to optimal data layout, particularly for unstructured grid methods. In the domain decomposition of implicit and inherently serial algorithms (e.g., the grid “sweep” in particle transport), parallel algorithms that maintain close-to-serial algorithm convergence rates are highly desirable. For both transport and hydrodynamics algorithms, efficient use of threads with shared memory is also of high interest.

Research Areas

- Adaptive algorithms
- Coupling to sub-grid scale models for material mixing and turbulence
- Local refinement and adaptation in space and angle in the areas of radiation and particle transport
- Algorithms that are highly efficient on parallel, cache-based computers

Computational Mechanics

ASAP is interested in research in the area of computational mechanics that will help to fully exploit scalable, terascale, massively parallel engineering mechanics simulations in support of the design and certification of nuclear weapon systems and components.

Research Areas

- Efficient and scalable solution strategies for nonlinear implicit transient dynamics
- Solution strategies for computing large deformations and fracture
- Improved contact algorithms
- Scalable view factor algorithms
- Scalable solution strategies for compressible and incompressible viscous flow
- Efficient algorithms for integrating complex nonlinear constitutive models in quasi-static applications and implicit dynamics applications involving large time steps
- Algorithms/models for energy dissipation at joints and interfaces in structural dynamics simulations

SELECTION

The White Paper should address how the proposed research meets the following factors. Selection of finalists for receipt of an RFP will be based on these factors.

- How the approach contributes to ASCI goals and technologies of importance to the DP Laboratories.
- How the university plans to integrate its work with the needs of the DP Laboratories in support of long-term ASCI goals.
- The impact the proposed research will have on the utilization of, and performance on, systems with thousands of processors and how that impact will be demonstrated.
- The degree to which the proposed research has value to more than one of the DP Laboratories.
- The level of focus on, and demonstration of, scalability to terascale computing.
- The extent to which the proposed work relates to the research issues outlined above.
- The proposed level and type of interaction between the Offeror's personnel and the DP Laboratories. (The purpose of such interactions is to facilitate technology and information exchange in both directions, as appropriate.)

For the current 13 Level 2 subcontractors, selection will also depend upon the quality of performance to date, the technical scope of proposed work, and the perceived value of that work relative to other White Papers received. The ASAP makes no guarantee either to award continuation subcontracts or to award subcontracts for new work. A significant portion of the annual funding (perhaps as much as 25%) may be made available for awards to new subcontractors based on the quality of the White Papers received. However, the ASAP reserves the right to allocate the available funding as it sees fit.

Each White Paper is limited to six pages in length and must address the following.

1. A description of the research effort.
2. Key personnel who will conduct and oversee the research.
3. How the research relates to the factors identified above.
4. How the participants will interact with the DP Labs.
5. A preliminary budget.
6. A preliminary period of performance.
7. Contacts: Name, address, e-mail address, and telephone and facsimile number of the university point of contact.

Address items 2, 6, and 7 on a single cover page; up to four pages may be devoted to items 1 and 3; and address items 4 and 5 using no more than 1.5 pages. A university may submit an additional vita per PI and/or Co-PI, limited to two pages in length and including the five most relevant references.

WHITE PAPER SUBMISSION REQUIREMENTS

Interested universities are asked to submit White Paper(s) as described above. White Paper(s) must be received at LLNL by **12:00 Midnight (PST), April 9, 2001** in order to receive consideration. Acceptance of late submissions will be at the sole discretion of the ASAP. White Papers will be accepted electronically in Microsoft Word (IBM or Macintosh compatible) or in PDF format. Submit electronic responses to Barbara Larson at larson4@llnl.gov. White Paper(s) not provided electronically must be submitted on 3.5-in. floppy disk or CD ROM along with one printed copy to the address below.

Address for Commercial Courier
Lawrence Livermore National Laboratory
Attention: Barbara Larson, L-550
7000 East Avenue
Livermore, CA 94550

Address for Mailing
Lawrence Livermore National Laboratory
Attention: Barbara Larson, L-550
P.O. Box 808
Livermore, CA 94551

QUESTIONS

Questions concerning this Request should be directed to Barbara Larson at (925) 422-0607 or e-mail at larson4@llnl.gov. In her absence, you may contact Brandt Esser at (925) 423-1518 or e-mail at esser3@llnl.gov. Address technical questions about specific research topics to the Points of Contact listed under the title of each topic.

PROPRIETARY INFORMATION

If proprietary data is included in a White Paper, the paper must be marked “Proprietary.” The DP Laboratories will maintain the proprietary data in confidence, giving it the same degree of care, but no less than a reasonable degree of care, as the DP Laboratories exercise with their own proprietary data to prevent its unauthorized disclosure.

Thank you for your interest. We look forward to your responses to this Request for Expressions of Interest.

Sincerely,

The ASCI Academic Strategic Alliances Program

This work was performed under the auspices of the University of California, Lawrence Livermore National Laboratory under Contract No. W-7405-Eng-48.